

Why Mapping the Internet is Hard

Aaron Clauset* and Cristopher Moore*,†

*Computer Science Department and †Department of Physics and Astronomy,
University of New Mexico, Albuquerque NM 87131

(aaron,moore)@cs.unm.edu

(Dated: February 2, 2008)

Despite great effort spent measuring topological features of large networks like the Internet, it was recently argued that sampling based on taking paths through the network (e.g., traceroutes) introduces a fundamental bias in the observed degree distribution. We examine this bias analytically and experimentally. For classic random graphs with mean degree c , we show analytically that traceroute sampling gives an observed degree distribution $P(k) \sim k^{-1}$ for $k \lesssim c$, even though the underlying degree distribution is Poisson. For graphs whose degree distributions have power-law tails $P(k) \sim k^{-\alpha}$, the accuracy of traceroute sampling is highly sensitive to the population of low-degree vertices. In particular, when the graph has a large excess (i.e., many more edges than vertices), traceroute sampling can significantly misestimate α .

The Internet is a canonical complex network, and a great deal of effort has been spent measuring its topology. However, unlike the Web, where the outgoing links are directly visible, we cannot typically ask a router who are its neighbors. As a result, studies have sought to *infer* the topology of the Internet by aggregating either paths through the network (i.e., traceroutes from a small number of sources to a large number of destinations) [1, 2, 3, 4, 5], routing decisions like those imbedded in BGP routing tables [6, 7, 8], or both [9, 10, 11]. Although such methods are known to be noisy [7, 12, 13], they strongly suggest that the Internet has a power-law degree distribution at both the router and domain levels.

However, Lakhina et al. [14] recently argued that traceroute-based sampling introduces a fundamental bias in topological inferences, since the probability that an edge appears within an efficient route decreases with the distance from the source. They showed empirically that traceroutes from a single source cause Erdős-Rényi random graphs $G(n, p)$, whose underlying distribution is Poisson [15], to appear to have a power law degree distribution $P(k) \sim k^{-1}$.

In this paper, we prove this result analytically by modeling the growth of a spanning tree on $G(n, p)$ using differential equations. Certainly no one would argue that the Internet is a purely random graph; indeed, the degree distributions reported in e.g. [9] have $P(k) \sim k^{-\alpha}$ with $2 < \alpha < 3$. However, it is evocative that traceroute sampling can create the appearance of a power-law degree distribution where none in fact exists.

Even if the Internet has a power-law degree distribution, it is reasonable to ask whether traceroute sampling gives an accurate estimate of the exponent α (a question raised also in [16]). Here, we demonstrate that power-law degree distributions are only well sampled when the graph has a small *excess*, i.e., a mean degree close to 2, so that the graph is very treelike. Other cases can result in a significant over- or under-estimation of α . Indeed, the accuracy of traceroute sampling is highly sensitive to the low-degree part of the degree distribution, not just the high-degree tail.

Traceroute spanning trees: analytical results. The set of traceroutes from a single source can be modeled as a spanning tree [17]. If we assume that Internet routing protocols approximate shortest paths, this spanning tree is built breadth-first from the source. In fact, the results of this section apply to spanning trees built in a variety of ways, as we will see below.

We can think of the spanning tree as built step-by-step by an algorithm that explores the graph. At each step, every vertex in the graph is labeled *reached*, *pending*, or *unknown*. Pending vertices are the leaves of the current tree; reached vertices are interior vertices; and unknown vertices are those not yet connected. We initialize the process by labeling the source vertex pending, and all other vertices unknown. Then the growth of the spanning tree is given by the following pseudocode:

```
while there are pending vertices:
    choose a pending vertex v
    label v reached
    for every unknown neighbor u of v,
        label u pending.
```

The type of spanning tree is determined by how we choose the pending vertex v . Storing vertices in a queue and taking them in FIFO (first-in, first-out) order gives a breadth-first tree of shortest paths; if we like we can break ties randomly between vertices of the same age in the queue, which is equivalent to adding a small noise term to the length of each edge as in [14]. Storing pending vertices on a stack and taking them in LIFO (last-in, first-out) order builds a depth-first tree. Finally, choosing v uniformly at random from the pending vertices gives a “random-first” tree.

Surprisingly, while these three processes build different trees, and traverse them in different orders, they all yield the same degree distribution when n is large. To illustrate this, Fig. 1 shows the degree distributions for each type of spanning tree for a random graph $G(n, p = c/n)$ where $n = 10^5$ and $c = 100$. The three degree distributions are indistinguishable, and all agree with the analytic results derived below.

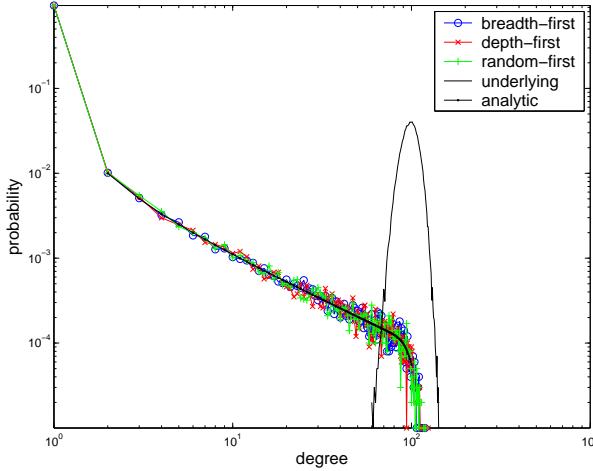


FIG. 1: Sampled degree distributions from breadth-first, depth-first and random-first spanning trees on a random graph of size $n = 10^5$ and average degree $c = 100$, and our analytic results (black dots). For comparison, the black line shows the Poisson degree distribution of the underlying graph. Note the power-law behavior of the apparent degree distribution $P(k) \sim k^{-1}$, which extends up to a cutoff at $k \sim c$.

We now show analytically that building spanning trees in Erdős-Rényi random graphs $G(n, p = c/n)$ using any of the processes described above gives rise to an apparent power law degree distribution $P(k) \sim k^{-1}$ for $k \lesssim c$. We focus here on the case where the average degree c is large, but constant with respect to n ; we believe our results also hold if c is a moderately growing function of n , such as $\log n$ or n^ϵ for small ϵ , but it seems more difficult to make our analysis rigorous in that case.

To model the progress of the `while` loop described above, let $S(T)$ and $U(T)$ denote the number of pending and unknown vertices at step T respectively. The expected changes in these variables at each step are

$$\begin{aligned} E[U(T+1) - U(T)] &= -pU(T) \\ E[S(T+1) - S(T)] &= pU(T) - 1 \end{aligned} \quad (1)$$

Here the $pU(T)$ terms come from the fact that a given unknown vertex u is connected to the chosen pending vertex v with probability p , in which case we change its label from unknown to pending; the -1 term comes from the fact that we also change v 's label from pending to reached. Moreover, these equations apply no matter how we choose v ; whether v is the “oldest” vertex (breadth-first), the “youngest” one (depth-first), or a random one (random-first). Since edges in $G(n, p)$ are independent, the events that v is connected to each unknown vertex u are independent and occur with probability p .

Writing $t = T/n$, $s(t) = S(tn)/n$ and $u(tn) = U(t)/n$, the difference equations (1) become the following system of differential equations,

$$\frac{du}{dt} = -cu, \quad \frac{ds}{dt} = cu - 1 \quad (2)$$

With the initial conditions $u(0) = 1$ and $s(0) = 0$, the solution to (2) is

$$u(t) = e^{-ct}, \quad s(t) = 1 - t - e^{-ct}. \quad (3)$$

The algorithm ends at the smallest positive root t_f of $s(t) = 0$; using Lambert's function W , defined as $W(x) = y$ where $ye^y = x$, we can write

$$t_f = 1 + \frac{1}{c}W(-ce^{-c}). \quad (4)$$

Note that t_f is the fraction of vertices which are reached at the end of the process, and this is simply the size of the giant component of $G(n, c/n)$.

Now, we wish to calculate the degree distribution $P(k)$ of this tree. The degree of each vertex v is the number of its previously unknown neighbors, plus one for the edge by which it became attached (except for the root). Now, if v is chosen at time t , in the limit $n \rightarrow \infty$ the probability it has k unknown neighbors is given by the Poisson distribution with mean $m = cu(t)$, $\text{Poisson}(m, k) = e^{-m}m^k/k!$. Averaging over all the vertices in the tree gives

$$P(k+1) = \frac{1}{t_f} \int_0^{t_f} dt \text{Poisson}(cu(t), k).$$

It is helpful to change the variable of integration to m . Since $m = ce^{-ct}$ we have $dm = -cm dt$, and

$$\begin{aligned} P(k+1) &= \frac{1}{t_f} \int_{c(1-t_f)}^c dm \frac{\text{Poisson}(m, k)}{cm} \\ &\approx \int_{ce^{-c}}^c dm \frac{\text{Poisson}(m, k)}{cm} \\ &= \frac{1}{ck!} \int_{ce^{-c}}^c dm e^{-m} m^{k-1}. \end{aligned} \quad (5)$$

Here in the second line we use the fact that $t_f \approx 1 - e^{-c}$ when c is large (i.e., the giant component encompasses almost all of the graph).

The integral in (5) is given by the difference between two incomplete Gamma functions. However, since the integrand is peaked at $m = k - 1$ and falls off exponentially for larger m , for $k \lesssim c$ it coincides almost exactly with the full Gamma function $\Gamma(k)$. Specifically, for any $c > 0$ we have

$$\int_0^{ce^{-c}} dm e^{-m} m^{k-1} < ce^{-c}$$

and, if $k - 1 = c(1 - \epsilon)$ for $\epsilon > 0$, then

$$\begin{aligned} \int_c^\infty dm e^{-m} m^{k-1} &= e^{-c} c^{k-1} \int_0^\infty dx e^{-x} (1 + x/c)^{k-1} \\ &< e^{-c} c^{k-1} \int_0^\infty dx e^{-x} e^{x(k-1)/c} \\ &= \frac{e^{-c} c^{k-1}}{\epsilon} < \frac{e^{-(k-1)} (k-1)^{k-1}}{\epsilon} \\ &\approx \frac{\Gamma(k)}{\epsilon \sqrt{2\pi(k-1)}} \end{aligned}$$

This is $o(\Gamma(k))$ if $\epsilon \gtrsim 1/\sqrt{k}$, i.e., if $k < c - c^\alpha$ for some $\alpha > 1/2$. In that case we have

$$P(k+1) = (1 - o(1)) \frac{\Gamma(k)}{ck!} \sim \frac{1}{ck} \quad (6)$$

giving a power law k^{-1} up to $k \sim c$.

Although we omit some technical details, this derivation can be made mathematically rigorous using results of Wormald [18], who showed that under fairly generic conditions, the state of discrete stochastic processes like this one is well-modeled by the corresponding rescaled differential equations. Specifically it can be shown that if we condition on the initial source vertex being in the giant component, then with high probability, for all t such that $0 < t < t_f$, $U(tn) = u(t)n + o(n)$ and $S(tn) = s(t)n + o(n)$. It follows that with high probability our calculations give the correct degree distribution of the spanning tree within $o(1)$.

Power-law degree distributions. We now turn to the case where the Internet *does* have a power-law degree distribution $P(k) \sim k^{-\alpha}$, and demonstrate that unless the *excess*, i.e., the number of edges minus the number of vertices, is small, traceroute sampling can significantly misestimate α .

There are several methods of constructing random graphs with power-law degree distributions and we use two to support our claim: the configuration model [19] in which the graph is random but conditioned on its degree distribution, and preferential attachment [20] in which the graph is grown by a dynamical process and has a degree distribution with a power-law tail.

In the configuration model, we examined graphs where $P(k) = k^{-\alpha}/\zeta(\alpha)$ for all $k \geq 1$, with α ranging from 1.5 to 3. Since these graphs are not necessarily connected, we compare the sampled and underlying degree distributions of the giant component (the latter has a power-law tail with the same exponent as the entire graph). Fig. 2 shows that as α increases, the observed distribution gets closer to the underlying distribution. This closeness is because, for instance, when $\alpha = 3$ the ratio of edges to vertices in the giant component is only 1.02 so its excess is only 0.02 per vertex. Thus *any* spanning tree on the giant component will include almost all of its edges, and sample its degree distribution fairly well.

However, the size and excess of the giant component are highly sensitive to the low-degree part of the degree distribution, not just its power-law tail. To illustrate this, Fig. 3 shows graphs grown using the preferential attachment model of [20]. Here every vertex has degree at least m , since it is given m edges at birth. As m increases, the slope of the observed distribution on a log-log plot becomes more shallow, giving a significant underestimate of α ; for instance, for $m = 4$ we observe a slope of 2.7 rather than the correct value $\alpha = 3$. Using the configuration model to construct random graphs with a minimum degree m and a degree distribution with a power-law tail yields similar results.

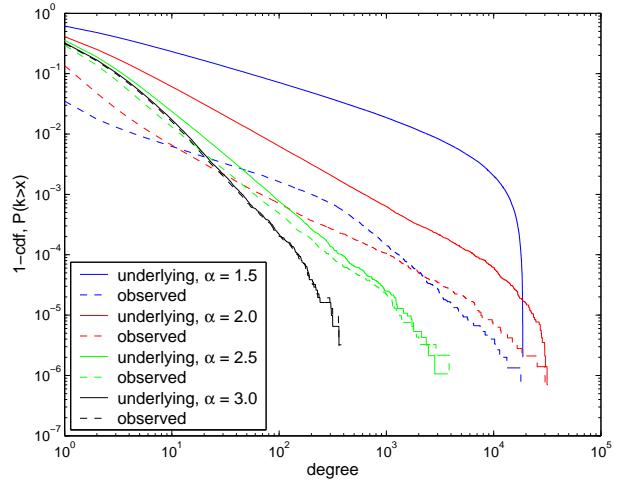


FIG. 2: Comparison of underlying and observed degree distributions in the configuration model with $n = 5 \times 10^5$ and various α .

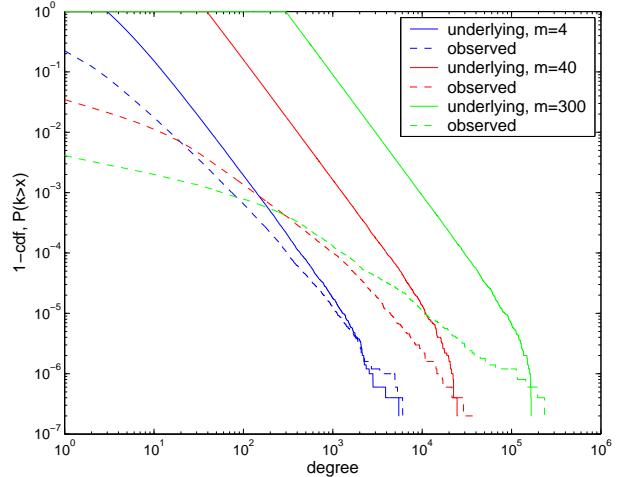


FIG. 3: Displacement of the power law tail for preferential attachment networks with $n = 5 \times 10^5$. Traceroute sampling significantly underestimates the slope α as m increases.

This underestimation of α occurs because traceroutes sample high-degree vertices more accurately than lower-degree ones: high-degree vertices are encountered early on in the breadth-first tree, when most of their neighbors are still unknown, while lower-degree vertices are encountered later, by which time most of their neighbors are already reached. Thus the “visibility” of a vertex’s edges increases with its degree [14], making the slope of the observed distribution less negative.

For smaller values of α , Fig. 2 shows that traceroute sampling encounters another kind of problem at smaller values of α , namely significant finite-size effects. The observed value of α is roughly correct up to a “knee,” above which the degree distribution falls off more sharply.

For $\alpha = 1.5$, for instance, Fig. 4 shows that this knee occurs at a degree $k \sim n^{0.5}$. In these cases, a linear fit to the observed degree distribution will considerably over-estimate α unless we omit the data above the knee.

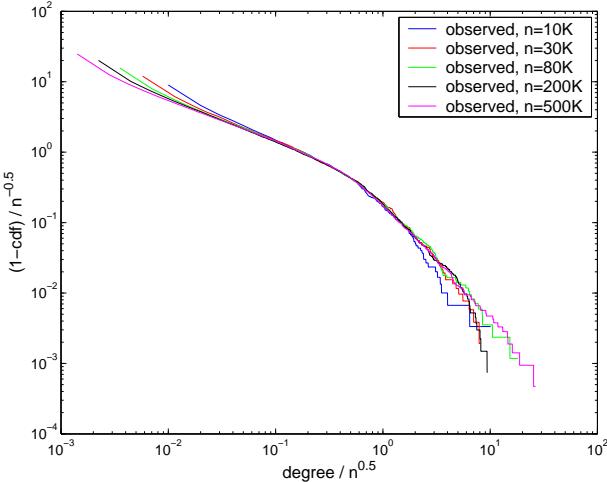


FIG. 4: Finite size effects for traceroute sampling with $\alpha = 1.5$, with data collapse for various n . Above a “knee” at $n^{0.5}$ the observed degree distribution falls off more sharply.

Conclusions. There are two properties of the Internet which make it difficult to map: unlike the World Wide Web where links are visible, the Internet’s topology must be queried indirectly, e.g., by traceroutes; and,

since efficient routing protocols cause these traceroutes to approximate shortest paths, edges far from the source are difficult to see. It was observed by [14] that these effects can significantly bias the observed degree distribution, and even create the appearance of a power law where none exists. We have proved this result analytically for random graphs $G(n, p)$, which yield an observed distribution $P(k) \sim k^{-1}$ for k up to the average degree. Other mechanisms by which power laws can appear in $G(n, p)$ include gradient-based flows [21], probabilistic pruning [22], and minimum spanning trees on weighted random graphs [23].

While it seems likely that the Internet does have a power-law distribution, we have shown that traceroute sampling can significantly misestimate the scaling exponent α . Thus we suggest that the published values of α may not accurately reflect the real scaling of the Internet’s topology. This poses an interesting inverse problem: namely, given the value of α observed in traceroutes, what is the most likely value of α in the underlying graph? Also, since traceroutes from a single source, or a small number of sources (briefly explored in [14, 16]), are inherently biased, how many sources are needed, as a function of network size and topology, to accurately sample the network?

We are grateful to David Kempe, Mark Newman, Mark Crovella, Paolo De Los Rios, Michel Morvan, Todd Underwood, Dimitris Achlioptas, Nick Hengartner and Tracy Conrad for helpful conversations. This work was funded by NSF grant PHY-0200909, and facilitated by Hewlett-Packard Gift No. 88425.1 to Darko Stefanovic.

- [1] J.-J. Pansiot and D. Grad, *ACM SIGCOMM Communication Review*, **28**(1):41-50, (1998)
- [2] R. Govindan and H. Tangmunarunkit, *Proc. IEEE INFOCOM* (2000).
- [3] CAIDA skitter Project. <http://www.caida.org/>
- [4] Opte Project. <http://www.opte.org>
- [5] Internet Mapping Project, Lumeta Corp.
<http://research.lumeta.com/ches/map/>
- [6] Autonomous Systems (ASs) use the Border Gateway Protocol (BGP) to route packets between networks. BGP routing tables represent logical domain adjacencies rather than physical connectivity. Tables contain advertised paths through this logical network; routing within domains is typically done via Open Shortest Path First or other distributed spanning tree protocol.
- [7] L. Amini, A. Shaikh and H. Schulzrinne, *Proc. SPIE IT-Com* (2002).
- [8] University of Oregon Route Views Project,
<http://www.antc.uoregon.edu/route-views>
- [9] M. Faloutsos, P. Faloutsos and C. Faloutsos, *Proc. ACM SIGCOMM* (1999) 251–62.
- [10] N. Spring, R. Mahajan and D. Wetherall, *Proc. ACM SIGCOMM* (2002).
- [11] Traceroute Organization. <http://www.traceroute.org>
- [12] Q. Chen, et al., preprint (2002).
- [13] P. Barford, A. Bestavros, J. Byers and M. Crovella, *SIGCOMM Internet Measurement Workshop* (2001).
- [14] A. Lakhina, J. Byers, M. Crovella and P. Xie, *Proc. IEEE INFOCOM* (2003).
- [15] P. Erdős and A. Rényi, *Publ. Math. Inst. Hung. Acad. Sci* 5 (1960) 17–61.
- [16] L. Dall’Asta, et al., preprint (2004).
- [17] While there are typically multiple shortest paths from the source to a given destination, their union forms a tree if we break ties e.g. by giving each node an index and taking the shortest path which comes first in lexicographic order.
- [18] N.C. Wormald, *Ann. Appl. Probab.* **5**(4) (1995) 1217–1235.
- [19] B. Bollobás, *Random graphs*. Academic Press, 1985.
- [20] A-L Barabási and R. Albert, *Science* **286**:509–512 (1999)
- [21] Z. Toroczkai and K.E. Bassler, *Nature* **428**:716 (2004)
- [22] T. Peterman and P. De Los Rios, *Eur. Phys. J. B*, to appear.
- [23] P.J. Macdonald, E. Almaas, and A.-L. Barabasi, preprint (2004).